9th USENIX Symposium on Operating Systems Design and Implementation (OSDI'10)

Vancouver, Canada
4 - 6 October 2010

ISBN: 978-1-7138-0463-5

Printed from e-media with permission by:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571

Some format issues inherent in the e-media version may also appear in this print version.

Copyright© (2010) by Usenix Association
All rights reserved.

Printed with permission by Curran Associates, Inc. (2020)

For permission requests, please contact Usenix Association
at the address below.

Usenix Association
2560 Ninth Street, Suite 215
Berkeley, California, 94710

https://www.usenix.org/

Additional copies of this publication are available from:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: 845-758-0400
Fax: 845-758-2633
Email: curran@proceedings.com
Web: www.proceedings.com

9th USENIX Symposium on Operating Systems Design and Implementation
October 4–6, 2010
Vancouver, BC, Canada

Message from the Program Co-Chairs. vii

Monday, October 4

Kernels: Past, Present, and Future

An Analysis of Linux Scalability to Many Cores 1
Silas Boyd-Wickizer, Austin T. Clements, Yandong Mao, Aleksey Pesterev, M. Frans Kaashoek, Robert Morris, and Nickolai Zeldovich, MIT CSAIL

Trust and Protection in the Illinois Browser Operating System 17
Shuo Tang, Haohui Mai, and Samuel T. King, University of Illinois at Urbana-Champaign

FlexSC: Flexible System Call Scheduling with Exception-Less System Calls 33
Livio Soares and Michael Stumm, University of Toronto

Inside the Data Center, 1

Finding a Needle in Haystack: Facebook’s Photo Storage. 47
Doug Beaver, Sanjeev Kumar, Harry C. Li, Jason Sobel, and Peter Vajgel, Facebook Inc.

Availability in Globally Distributed Storage Systems 61
Daniel Ford, François Labelle, Florentina I. Popovici, Murray Stokely, Van-Anh Truong, Luiz Barroso, Carrie Grimes, and Sean Quinlan, Google, Inc.

Nectar: Automatic Management of Data and Computation in Datacenters 75
Pradeep Kumar Gunda, Lenin Ravindranath, Chandramohan A. Thekkath, Yuan Yu, and Li Zhuang, Microsoft Research Silicon Valley

Security Technologies

Intrusion Recovery Using Selective Re-execution. 89
Taesoo Kim, Xi Wang, Nickolai Zeldovich, and M. Frans Kaashoek, MIT CSAIL

Static Checking of Dynamically-Varying Security Policies in Database-Backed Applications. 105
Adam Chlipala, Impredicative LLC

Accountable Virtual Machines. 119
Andreas Haeberlen, University of Pennsylvania; Paarijaat Aditya, Rodrigo Rodrigues, and Peter Druschel, Max Planck Institute for Software Systems (MPI-SWS)

Concurrency Bugs

Bypassing Races in Live Applications with Execution Filters 135
Jingyue Wu, Heming Cui, and Junfeng Yang, Columbia University

Effective Data-Race Detection for the Kernel 151
John Erickson, Madanlal Musuvathi, Sebastian Burckhardt, and Kirk Olynyk, Microsoft Research

Ad Hoc Synchronization Considered Harmful 163
Weiwei Xiong, University of Illinois at Urbana-Champaign; Soyeon Park, Jiaqi Zhang, and Yuanyuan Zhou, University of California, San Diego; Zhiqiang Ma, Intel

Tuesday, October 5

Deterministic Parallelism

Deterministic Process Groups in dOS 177
Tom Bergan, Nicholas Hunt, Luis Ceze, and Steven D. Gribble, University of Washington

Efficient System-Enforced Deterministic Parallelism 193
Amittai Aviram, Shu-Chun Weng, Sen Hu, and Bryan Ford, Yale University

Stable Deterministic Multithreading through Schedule Memoization 207
Heming Cui, Jingyue Wu, Chia-che Tsai,and Junfeng Yang, Columbia University

Systems Management

Enabling Configuration-Independent Automation by Non-Expert Users 223
Nate Kushman and Dina Katabi, Massachusetts Institute of Technology

Automating Configuration Troubleshooting with Dynamic Information Flow Analysis 237
Mona Attariyan and Jason Flinn, University of Michigan

Inside the Data Center, 2

Large-scale Incremental Processing Using Distributed Transactions and Notifications 251
Daniel Peng and Frank Dabek, Google, Inc.

Reining in the Outliers in Map-Reduce Clusters using Mantri 265
Ganesh Ananthanarayanan, Microsoft Research and UC Berkeley; Srikanth Kandula and Albert Greenberg, Microsoft Research; Ion Stoica, UC Berkeley; Yi Lu, Microsoft Research; Bikas Saha and Edward Harris, Microsoft Bing

Transactional Consistency and Automatic Management in an Application Data Cache 279
Dan R.K. Ports, Austin T. Clements, Irene Zhang, Samuel Madden, and Barbara Liskov, MIT CSAIL

Piccolo: Building Fast, Distributed Programs with Partitioned Tables 293
Russell Power and Jinyang Li, New York University

Cloud Storage

Depot: Cloud Storage with Minimal Trust. 307
Prince Mahajan, Srinath Setty, Sangmin Lee, Allen Clement, Lorenzo Alvisi, Mike Dahlin, and Michael Walfish, The University of Texas at Austin

Comet: An Active Distributed Key-Value Store 323
Roxana Geambasu, Amit A. Levy, Tadayoshi Kohno, Arvind Krishnamurthy, and Henry M. Levy, University of Washington

SPORC: Group Collaboration using Untrusted Cloud Resources 337
Ariel J. Feldman, William P. Zeller, Michael J. Freedman, and Edward W. Felten, Princeton University

Wednesday, October 6

Production Networks

Onix: A Distributed Control Platform for Large-scale Production Networks. 351
Teemu Koponen, Martin Casado, Natasha Gude, and Jeremy Stribling, Nicira Networks; Leon Poutievski, Min Zhu, and Rajiv Ramanathan, Google; Yuichiro Iwata, Hiroaki Inoue, and Takayuki Hama, NEC; Scott Shenker, International Computer Science Institute (ICSI) and UC Berkeley

Can the Production Network Be the Testbed? 365
Rob Sherwood, Deutsche Telekom Inc. R&D Lab; Glen Gibb and Kok-Kiong Yap,Stanford University; Guido Appenzeller, Big Switch Networks; Martin Casado, Nicira Networks; Nick McKeown and Guru Parulkar, Stanford University

Building Extensible Networks with Rule-Based Forwarding 379
Lucian Popa, University of California, Berkeley, and ICSI, Berkeley; Norbert Egi, Lancaster University; Sylvia Ratnasamy, Intel Labs, Berkeley; Ion Stoica, University of California, Berkeley

Mobility

TaintDroid: An Information-Flow Tracking System for Realtime Privacy Monitoring on Smartphones 393
William Enck, The Pennsylvania State University; Peter Gilbert, Duke University; Byung-gon Chun, Intel Labs; Landon P. Cox, Duke University; Jaeyeon Jung, Intel Labs; Patrick McDaniel, The Pennsylvania State University; Anmol N. Sheth, Intel Labs

StarTrack Next Generation: A Scalable Infrastructure for Track-Based Applications 409
Maya Haridasan, Iqbal Mohomed, Doug Terry, Chandramohan A. Thekkath, and Li Zhang, Microsoft Research Silicon Valley

Virtualization

The Turtles Project: Design and Implementation of Nested Virtualization 423
Muli Ben-Yehuda, IBM Research—Haifa; Michael D. Day, IBM Linux Technology Center; Zvi Dubitzky, Michael Factor, Nadav Har'El, and Abel Gordon, IBM Research—Haifa; Anthony Liguori, IBM Linux Technology Center; Orit Wasserman and Ben-Ami Yassour, IBM Research—Haifa

mClock: Handling Throughput Variability for Hypervisor IO Scheduling. 437
Ajay Gulati, VMware Inc.; Arif Merchant, HP Labs; Peter J. Varman, Rice University

Virtualize Everything but Time 451
Timothy Broomhead, Laurence Cremean, Julien Ridoux, and Darryl Veitch, Center for Ultra-Broadband Information Networks (CUBIN), The University of Melbourne